Faculty Research Working Papers Series

Parallel R&D Paths Revisited

F. M. Scherer
John F. Kennedy School of Government – Harvard University

September 2007

RWP07-040

PARALLEL R&D PATHS REVISITED

F. M. Scherer
Harvard University
Revised September 2007

1. Introduction

This paper revisits the role pursuing parallel paths -- i.e., supporting simultaneously a diversity of experiments or product designs to hedge against uncertainties in securing a desired technological result -- plays in research and development strategy. Some of the more novel points advanced here come from my early publications and from unpublished lectures in a course on the economics of technological innovation and economic growth, taught repeatedly between 1982 and 2005. Some are newer, worked out to fill lacunae I had left in past treatments.

My introduction to the parallel paths strategy came in joint research with M. J. Peck on advanced weapons research and development. In our joint book,[1] Peck and I proposed, expanding upon a suggestion we heard first at Bell Telephone Laboratories, that the scheduling of an R&D project entailed a tradeoff between speed of development and cost. Figure 1 reproduces our diagram introducing the basic concept. One could accelerate the expected completion date of a development project, but (within efficient tradeoff curve segment NR) only by incurring higher cost. Time could be saved by assigning more talent to the effort, but only subject to diminishing marginal returns; by overlapping tests before the first step has yielded all the information useful in a later step; and by hedging against uncertainty by supporting parallel research or development approaches. The essence of the parallel paths strategy was:[2]

> ... operating simultaneously two or more approaches to the step, test, or problem to insure that at least one approach will hit the mark at the earliest possible moment.

We proposed a crude calculus-based solution to the time-cost tradeoff problem,

---

1 . M. J. Peck and F. M. Scherer, The Weapons Acquisition Process: An Economic Analysis (Harvard Business School Division of Research: 1962).

2 . Ibid., p. 261.

finding an optimum where the (negative) slope of the convex time-cost tradeoff function was equal to the time derivative of a function measuring the expected military value from successful completion of the research and development project. The more valuable the completed weapon system, the more rapidly it should be developed, among other things through the use of parallel paths.

There is no systematic evidence on how widely parallel paths strategies are used, but there is reason to believe the best-known cases are not atypical. Thomas Edison is said to have tested 1,600 materials for his electric lamp filaments before focusing on a carbonized design. In 1934 DuPont synthesized 81 different polyamide compounds in its quest for what eventually became nylon. Five were carried into further experiments.[3] A pioneer in science-guided rational drug design explored 367 different molecules before finding one with good prospects for suppressing the human body's rejection of artificial organ transplants.[4] Peck and I observed that the military authorities authorized parallel paths, sometimes in head-to-head competitions and sometimes more informally, in their quest to develop new fighters, bombers, and guided missiles. More recently, several design-stage alternatives were supported, and two full-scale prototypes were built for competitive evaluation during the 1990s, as precursors to what eventually became the F-22 "Raptor" advanced tactical fighter and the F-35 joint strike fighter. In perhaps the most famous case of all, U.S. defense authorities initially supported five different approaches to the problem of producing fissionable material for an atomic bomb, each expected in May 1942 to cost approximately $100 million, and four were sustained into production during the atomic bomb development effort.[5] And two

---

3 . David A. Hounshell and John Kelly Smith, Science and Corporate Strategy (Simon & Schuster: 1986), p. 259.

4 . Barry Werth, The Billion Dollar Molecule (Simon & Schuster: 1994), p. 251.

5 . See e.g. R.G. Hewlett and O.E. Anderson, The New World (vol. I of an Atomic Energy Commission history; Pennsylvania State University Press: 1962); and Richard Rhodes, The Making of the Atomic Bomb (Simon & Schuster: 1986). In The Economics of Defense in the Nuclear Age (Harvard University Press: 1960), p. 249, Charles J. Hitch and Roland McKean write that "... the method that succeeded in producing the material for the first bomb was regarded at first as among the least promising..." They do not identify the successful method or pinpoint the timing. In fact, the five methods' perceived prospects changed ranks several times between 1940 and 1943. In December 1941, the best alternatives were considered to be gaseous diffusion, which was combined with electromagnetic separation to produce material for the Hiroshima bomb, and centrifugal separation, which proved intractable and whose production plant authorization was cancelled in November 1942. It later became low-income nations'

different bomb designs --a gun-barrel design using uranium and an implosion device using plutonium -- were supported to the end of the War and afterward, the former exploded over Hiroshima and the latter at Nagasaki.

After completing my work with Peck, I sought to extend insights from the time-cost tradeoff concept, initially making my mathematical treatment more general and more rigorous in the context of government-supported R&D[6] and then building models showing how and when competition accelerated the pace of innovation in the civilian sector.[7]

The time-cost tradeoff approach was not without critics. At an informal meeting in the early 1960s, General Bernard Schriever, head of the U.S. Air Force's ballistic missile development program, insisted that there was no tradeoff: the quickest approach was also the least expensive. His implicit emphasis was on what we called the overhead effect, characterized by segment MN in Figure 1. There were also rumblings of skepticism from another important West Coast institution, the RAND Corporation. After my early work on time-cost tradeoffs was completed, economist Thomas Marschak published a rich set of what might be called impossibility theorems suggesting that parallel paths strategies could lead to an inverse convex relationship between development time and cost, but that important exceptions could also exist.[8]

2. Richard Nelson's Contribution

A more focused and widely-disseminated contribution came from Richard R.

---

preferred method.

6 . F. M. Scherer, "Government Research and Development Programs," in Robert Dorfman, ed., Measuring Benefits of Government Investments (Brookings Institution: 1965), pp. 12-57.

7 . See especially F. M. Scherer, "Research and Development Resource Allocation under Rivalry," Quarterly Journal of Economics (August 1967), pp. 359-394.

8 . Thomas Marschak, "Toward a Normative Theory of Development," in Marschak et al., eds., Strategy for R&D: Studies in the Microeconomics of Development (Springer-Verlag, 1967). For later, more specific, modelling approaches, see William J. Abernathy and Richard Rosenbloom, "Parallel and Sequential R&D Strategies," IEEE Transactions on Engineering Management (March 1968); and "Parallel Strategies in Development Projects," Management Science (June 1969), pp. 486-505.

Nelson,[9] who had been a colleague of Marschak at RAND. Nelson emphasized the beneficial role of parallel approaches in the face of R&D uncertainties, concluding (p. 363) that "we should be wary in damning the wastefulness of independent and competitive efforts" and more generally that "the number of alternative inventors ... should be greater, the greater the demand for the invention."

Nelson motivates his more general model with a simple numerical illustration. The objective is to develop a successful new fighter aircraft. At the outset, it is uncertain which of various alternative designs is likely to be successful. A prediction of future success is obtained by building and test-flying one or more prototypes embodying proposed designs. Each prototype effort costs $10 million (an astounding but plausible number, compared to present-day fighter aircraft development program costs in the billions of dollars!) and takes 20 months. Prototype tests can reveal any given design to be a more desirable Type I, with an expected further development cost to successful completion of $50 million over an additional time span of 20 months, or a less desirable Type II, for which the additional development effort is expected to cost $100 million and take 50 additional months. The a priori probability of a Type I outcome is 0.40 and for Type II 0.60. If only one design proves after prototype tests to be of Type I, it is carried into final development. If parallel prototype paths are pursued and more than one design proves to be a Type I, one of the successes is chosen randomly for further development. If none is a Type I, one of the Type II prototypes is selected randomly for high-cost final development.

Weighting outcomes by probabilities, Nelson computes the expected values of total development cost and time for alternative strategies -- supporting only a single first-stage prototype, chosen at random, or pursuing from two to five parallel prototype development paths. The outcomes, given his assumptions, are shown in Figure 2. Authorizing a single prototype leads to an expected total cost of $90 million and a probability-weighted time to completion of 58 months; with two approaches in parallel, development is not only faster -- 50.8 months -- but less expensive ($88 million). Choosing between one and two paths, there is no tradeoff: the two-path strategy is dominant. For the specific values chosen, a partial counter-

---

9 . Richard R. Nelson, "Uncertainty, Learning, and the Economics of Parallel Research and Development," Review of Economics and Statistics (November 1961), pp. 351-368. Peck and I were not aware of Nelson's article at the time we completed our Weapons Acquisition Process manuscript in the fall of 1961 (and Peck had taken a position in the Office of the Secretary of Defense).

example to my assumption of tradeoff curve convexity is demonstrated. Nelson's initially articulated criterion is "achieving a given objective at minimum cost," so it would appear that the two-path strategy is optimal. However, he recognizes that delay can also be costly. For strategies involving more than two parallel paths, time is reduced, but only at higher expected cost. A tradeoff materializes. Assuming (p. 360) that a month's delay in effect costs $1 million (i.e., that the Air Force is willing to spend an extra $1 million for each month saved), he finds the least-cost strategy to be pursuing three parallel paths.

Nelson's pioneering analysis makes a compelling case for the potential attractiveness of parallel R&D paths strategies. There are, however, two noteworthy problems.

First, his "cost of delay" or (in footnote 14) indifference curve relating delay cost to time[10] overlooks a rather general point. R&D is normally an investment made for the purpose of securing future benefits, which, as in most investment problems, accrue over a period of time. When the project is completed, one taps into a stream of benefits. What one loses by taking longer to complete one's development is the benefit from that stream during the period of delay. As Peck and I formulated crudely (smoothing unique and short-duration but uncertain combat needs with the probability of a combat situation) and I proposed more rigorously in 1965, the problem of optimal development timing consists of maximizing the difference between discounted benefits and R&D costs, i.e.,

$$(1) \quad \underset{T}{\text{Max}} \int_{T}^{H} v(t)\, e^{-rt}\, dt - C(T);$$

where v(t) is the depth of the benefits stream at time t, r is a conventional time discount rate (or in business problems, the so-called "hurdle rate"), C(T) is a convex inverse function relating development cost to the expected time for development completion, T is the time when the development is completed and benefits begin flowing in, and H is the decision-maker's time horizon. To be sure, erratic benefit stream configurations might require the single-valued time cost assumption of Nelson and Marschak, but equation (1) is more general and more consistent with the accepted literature on capital investment.

---

10 . See also Marschak, pp. 207-210, who also views delay as a cost and postulates indifference curves to resolve tradeoffs.

Second, Nelson's analysis compares only options entailing multiple but simultaneous prototype paths against the one-path alternative. He ignores alternative series scheduling strategies. Given his assumptions, a series strategy would build and test one prototype, determine at the end of 20 months whether it is a Type I or Type II, commence full-scale development if it is a Type I, and (here is the difference) shut down the first project and begin a second prototype project if the first prototype is found to be a Type II. The same decision rule could be followed when test results are obtained from the second prototype, and so on for as many iterations in series as one wishes to entertain. Where 0.6 is the probability of a Type II outcome for a single prototype and 0.4 the probability of a Type I outcome, the expected cost of completion for a three-stage series strategy (with full-scale development proceeding regardless of the third-stage results, if it is reached) is:

(2) $E(T) = \$10 \text{ million} + (0.4) \times \$50 \text{ million}$ [1st stage success]
$+ (0.6 \times \$10 \text{ million})$ [2nd prototype after failure]
$+ (0.6 \times 0.4) \times \$50 \text{ million}$ [2nd stage success)
$+ (0.6^2 \times \$10 \text{ million})$ [3rd proto. after 2nd failure]
$+ (0.6^2 \times 0.4) \times \$50 \text{ million}$ [3rd stage success]
$+ (0.6^3 \times \$100) \text{ million}$ [3rd stage failure]

$= \$80.4 \text{ million}.$

By a similar probability-weighted calculation, one finds that the expected time to completion with a three-stage strategy is 65.7 months. Although the expected time to completion is longer, the expected cost with the three-stage series strategy is lower than with any of the parallel paths strategies. The tradeoff is restored.

Figure 3 adds to Figure 2 three series time-cost tradeoff outcomes, for two-stage, four-stage, and six-stage strategies. Again, one sees that the tradeoff is restored, the only anomaly being the higher cost with a one-stage strategy than with two parallel paths. Whether one would choose to use a series strategy instead of a parallel paths strategy depends upon the depth of the benefits stream tapped when the R&D project is successful. The deeper the benefits stream, the more the optimum moves to the northwest into multiple parallel paths. For shallow expected benefits streams, series scheduling could be optimal. Small-scale research in a university setting characteristically emphasizes the series approach. However, for important problems -- those whose solution will tap deep benefits streams -- multiple investigators are likely to be working in parallel. Thus, real-world behavior --to be sure, not necessarily optimal -- may combine series and parallel strategies.

3. An Intermediate Step

To set the stage for my work on how rivalry in the private sector affects the speed of innovation, I considered it essential first to reaffirm that there was indeed a tradeoff between development time and cost, especially in the context of parallel path vs. series strategies under uncertainty. Several approaches to the problem were analyzed and published in a 1966 paper.[11] That paper verified a robust time-cost tradeoff under uncertainty. Having resolved that question to my satisfaction, essentially at the sub-optimizing research strategy level, I was able to simplify my analysis of optimal R&D strategies under rivalry by assuming the time-cost tradeoff function to be deterministic. The analysis of time-cost tradeoffs under uncertainty was published in what I now know to be a journal virtually unknown to economists and therefore was largely ignored, even though it was referenced in the first footnote to my more ambitious 1967 article.[12] For example, in her authoritative survey article on the economics of innovation timing, Jennifer Reinganum identifies two basic paradigms, one deterministic and one stochastic, and places my work only in the deterministic category, overlooking and not citing my earlier paper on tradeoffs under uncertainty. In hindsight, I realize it was a mistake for me not to submit my 1966 paper to the Review of Economics and Statistics as a follow-on to Richard Nelson's pioneering contribution. I did not do so because, as a newly-minted Ph.D., I was ignorant of how the publication process worked in economics. Also, at the time there was very little interest in the economics profession on the microeconomics of technological innovation. Since my 1966 paper appeared to be well outside the main stream, I sought the advice of a more experienced colleague and followed it, submitting the paper to the Naval Research Logistics Quarterly.

The 1966 paper established several propositions, summarized as follows:

(1) When each alternative research project has the same probability of success and the same cost, and when, for a series strategy, equal numbers of projects are scheduled per time period, a convex time-cost tradeoff function exists.

---

11 . "Time-Cost Tradeoffs in Uncertain Empirical Research Projects," Naval Research Logistics Quarterly, vol. 13 (March 1966), pp. 71-82. Material inadvertently omitted from the original version was published in the September 1966 edition of the same journal.

12 . Supra note 4.

(2) Under the assumptions of (1), the equal projects per time period strategy was not optimal, although the time-cost tradeoff continued to exist. Rather, costs were reduced, all else equal, by scheduling relatively few projects in the first period and then increasing the number of projects progressively in later periods. In a dynamic programming example with an equal individual project success probability of 0.05 and a cumulative success probability target of 0.95, the optimal number of projects in six successive periods was 6, 7, 8, 10, 12, and 17. However, costs were not highly sensitive to modest deviations from this optimal pattern.

(3) When individual project success probabilities and/or costs differ, one schedules first in a parallel-series strategy the projects with the highest success probabilities. The negative time-cost tradeoff persists. Letting project costs differ too complicated the analysis beyond the bounds of known computational feasibility.

(4) When completing one or more projects generates information that increases the success prospects of subsequent projects, the case for series scheduling is strengthened without eliminating the existence of a time-cost tradeoff.

(5) Simultaneous cross-project learning (cross fertilization) or other scope economies strengthens the case for parallel paths and might reverse the time-cost tradeoff for excessively serial strategies.

(6) All of the above analyses assumed that successful trials are near-perfect substitutes of approximately equal utility, so in a series strategy, testing ceases when one success has been achieved. When outcomes are differentiated so that additional outcomes have value, the case for series scheduling is weakened but not eliminated and the case for parallel scheduling is strengthened. We return to this piece of unfinished business in a later section.

4. Finding the Global Optimum

All of these analyses were focused narrowly on testing, and in most cases supporting, the existence of a negatively sloped time-cost tradeoff function under significant uncertainty. Once that function is established, the problem remains of finding the time-cost combination that maximizes the expected surplus of benefits over R&D costs. A simple but fairly general analysis affirmed that the deeper the stream of benefits tapped following successful R&D project completion, the more one optimally emphasized saving time over saving cost. To achieve insight into how sensitive net profits were to the pursuit of parallel and series strategies, a quantitative analysis of diverse success probability, benefit stream depths, and scheduling strategies was conducted. Where t was a running time variable, b_t was

the dollar value of the benefits realizable in the t^{th} time period contingent upon success, *M* was the cost per research approach, *q* was the probability that any given approach would fail (like *M*, assumed constant), *N* was the number of approaches originally scheduled (subject to reduction if an early success emerged), and *r* was the discount rate, the objective was to maximize net expected present value *V*, defined as:

$$(3)\quad V = \sum_{t=2}^{T} [1 - q^{N(t-1)/T}] [1 / (1+r)] b_t + \sum_{t=T+1}^{H} [1 - q^{N}] [1 / (1+r)_t b_t - \sum_{t=1}^{T} [q^{N(t-1)/T}] [1 / (1+r)^{t-1}] (NM/T)$$

with respect to N and T. The first term is the benefits received contingent upon early success in the planned experimentation period, the second the benefits after completed experimentation has (with some probability) yielded success, and the third the discounted present value of R&D costs.

Figure 4 reproduces the main results of the numerical search for optimal solutions, where each experiment is conducted within a single year, the cost per experiment is $1,000, potential benefits are measured in thousands of dollars per year, they continue out to year 25, and the time discount rate *r* is 0.06.[13] Verifying prior insights, one found that the optimal number of independent research projects increases monotonically with the depth of the benefits stream b_t. The article reported that *V* was relatively insensitive to the diverse combinations of *N* (the total number of planned experiments, if early success were not achieved) and *T* (the number of periods over which experimentation might continue). In other words, getting the total number of scheduled experiments right was much more important than the way they were scheduled over time. But getting *N/T* right was important, as Figure 4 shows. Finally, changes in the net benefit-maximizing *N/T* were more

13 . Benefits were assumed to begin flowing in at the earliest in year 2 and to be discounted at the end of the year in which they were realized.

sensitive to the depth of the benefit stream, the lower the probability of success in a single experiment was -- i.e., the greater the uncertainty.

As I reconsidered the relevance of these results to an analysis of uncertainty-hedging in pharmaceutical R&D,[14] I realized that I wanted to know more about the relationship between *N* and *T*, that is, on the extent of reliance upon series as compared to parallel scheduling of individual projects. Unfortunately, the ravages of time had obliterated my individual computations and the detailed assumptions underlying them. I therefore replicated the analysis anew for a plausible array of scheduling assumptions. The computations were done mainly for the case in which scheduling was expected to be most sensitive to differences in benefit stream depths, i.e., with a low 0.01 probability of success. This is akin to conditions in pre-clinical animal model tests to discover therapeutically interesting pharmaceutical molecules before testing in humans begins. Profit-maximizing solutions were found by inspection from profit vectors for alternative numbers of trials per period, given differing benefit stream depths (a relatively easy task in the era of spreadsheets). For the series strategies, the number of trials scheduled per period was equal, contrary to insight (2) from my 1966 tradeoff paper.

Figure 5 shows that the discounted profit-maximizing number of trials scheduled per period (*N/T)* differs widely, depending upon whether all tests are scheduled for the first period as compared to being spread conditionally over two, three, or four successive periods. When everything is done in the first period, by far the largest number of trials in a period is scheduled. The more periods over which the trials are spread, the smaller is the profit-maximizing number of trials per period -- in the hope that an early success will alleviate the need for later trials.[15] The optimal total number of projected trials increases two-to-threefold, however, as one moves from a one-stage to a four-stage strategy. If one is unlucky and no success is achieved in early stages, the larger number of trials will cost more than does the single-stage strategy. But that multi-stage cost is reduced by the probability that success will be achieved earlier and the later trials will be unnecessary.

---

14 . F. M. Scherer, "Pharmaceutical Innovation," to appear in Bronwyn Hall and Nathan Rosenberg, eds., Handbook on the Economics of Technical Change.

15 . The curves are discontinuous below b_t of 10, since no parallel or series strategies yielded positive net profits for benefits of $7,500 per year -- the next lowest value for which optima were calculated -- or less.

Figure 6 tests for the sensitivity of total discounted profits, i.e., expected benefits minus expected R&D costs, to alternative series scheduling strategies. Spreading trials, whose number is optimized for the series strategy chosen, over two periods is substantially more profitable in the net than running all trials simultaneously in the first period. This is so even in the absence of learning from unsuccessful tests, as assumed throughout this computation. However, profits are not very sensitive to moving from two to three or four trial periods. Thus, a bit of series scheduling appears to be a good thing, but diminishing returns set in rapidly.

To be sure, a disadvantage of spreading any given number of trials (actually variable among series alternatives in the Figure 6 computation) over more periods is a longer expected period of development. For $b_t = 25$, the expected development time, given the profit-maximizing choice of trial numbers, varies with the number of periods over which the trials are spread (i.e., with increasing recourse to series scheduling) as follows:

	E(T)
All in first period	1.00 years
Two periods	1.42 years
Three periods	1.70 years
Four periods	1.92 years

Since series scheduling tends to be less costly, all else equal, this implies again the existence of a time-cost tradeoff. In the computation conducted, this tradeoff is taken into account explicitly by choosing trial numbers that maximize net profits, i.e., discounted benefits less discounted R&D costs, compensating for waiting longer on average with more protracted series strategies to tap the benefits stream.[16]

16 . There is also a third variable. Since costs are lower with series strategies, all else equal, more trials are conducted when early trials yield no successes, and as a result, at the end of the sequence, the cumulative probability of success is higher. For $b_t = 25$, the total number of trials in the worst case is 108 for a fully parallel strategy (T = 1), 174 for two stages, 222 for three stages, and 264 for four stages. The cumulative success probabilities are correspondingly 0.662, 0.826, 0.893, and 0.930. A higher cumulative success probability, like lower trial costs, enhances net profits, which is also taken into account in the computations.

Because calculating the optimal number of trials is fraught with estimation uncertainties in real-world practice, Figure 7 tests for sensitivity to a crude second-best strategy: conducting under any of four different series assumptions the number of trials per period optimal for a two-period strategy.[17] Again, the profit difference between a one-period and two-period strategy is substantial. But for a larger number of periods, the profit sacrifice from using this second-best strategy is even smaller than under the assumptions of Figure 6.

Clearly, impressive profits are realized in most of the cases analyzed, even though there is substantial duplication of R&D costs. They rise nonlinearly, needless to say, with the depth of the annual benefits stream. For perspective, the discounted present value (at 6 percent simple end-of-year interest) of benefits starting in year 3 and ending in year 25 is 273.75 (thousands of dollars) with annual benefits of 25 (thousands) per year and 547 with benefits of 50 per year. Thus, for the least profitable single-stage strategy summarized by Figure 6, net profits (after the deduction of R&D costs) are 32 percent of maximum attainable benefits in the b_t = 25 case and 58 percent in the b_t = 50 case.[18]

Figure 8 broadens the perspective to show net profits from fully parallel (single-year) R&D strategies for a broader array of success probabilities. Quite plausibly, one observes much higher net profits from parallel paths strategies with single-trial success probabilities greater than our initially assumed 0.01, with commensurately smaller optimal trial numbers. The profit increases as success probabilities are raised from 0.05 to 0.20 are considerably smaller than those for increases from 0.01 to 0.05. At b_t = 25, net profits with optimal fully parallel paths and a single-trial success probability of 0.20 are 93 percent of maximum attainable benefits, calculated as in the previous paragraph. One sees too that parallel paths strategies yield positive net profits for benefit stream depths considerably lower than 10 (thousand) per year, which was the approximate breakeven threshold with a single-trial success probability of only 0.01.

5. Multiple and Diverse Payoff Cases

---

17 . For the one, three, and four-stage strategies, these are not the profit-maximizing strategies.

18 . Net profits as a percent of actual discounted benefits, the latter reduced by less than unit probabilities of ultimate success, are necessarily smaller. See also Scherer, "Pharma-ceutical Innovation," supra note 14.

The analyses presented thus far have assumed consistently that there is uncertainty as to which research project or task will yield a good solution, but once a single solution is found, it suffices and the investigation can end. In effect, the main uncertainty is scientific or technological. But as recognized in generalization (6) above, a parallel paths investigation may yield more than one good solution. Consumer tastes differ, and one solution may satisfy one set of consumer wants while others better meet other consumers' wants. Moreover, there is abundant evidence that some research and development results, though technically successful, elicit relatively little consumer demand while others turn out to be "blockbusters." The uncertainties here are more on the demand side than the technological side. Quantitative research has shown conclusively that the size distribution of profits from technological innovations is highly skew.[19] That is, most R&D projects achieving technologically successful results yield only modest profits, but the distribution has a very long right-hand tail encompassing very high profit returns. On average, the top ten percent of R&D project outcomes by number yield roughly 50 to 80 percent of total portfolio profits, depending in part upon the R&D stage (i.e., concept-discovery vs. final development) whose outcomes are measured.

The diversity of consumer preferences combines with uncertainty as to which solutions best satisfy particular consumer clusters' wants to make investing in R&D like throwing darts with imprecise aim at a dartboard, over which payoffs are distributed randomly but with a skew distribution. To illustrate this proposition, a Monte Carlo experiment was conducted.[20] One hundred possible R&D outcomes were defined, with each of which a randomized discounted quasi-rent payoff was associated. Consistent with the empirical evidence, the distribution followed a log normal distribution of the form $10^{\text{normal}(0,1)} \times 1000$. For the payoff matrix (i.e., dartboard) used in all iterations of the experiment, the mean payoff was $5,490, the median $1,310, and the maximum payoff $91,700. The top ten payoffs accounted for 63.1 percent of the total across all 100 payoffs.

If innovative projects could be targeted precisely toward the most lucrative products or processes, innovators would in effect direct darts toward all prospects with expected quasi-rents exceeding R&D costs. Assuming a uniform R&D cost of

---

19 . See F. M. Scherer and Dietmar Harhoff, "Technology Policy for a World of Skew-Distributed Outcomes," Research Policy, vol. 29 (April 2000), pp. 559-566.

20 . This section is taken with modifications from F. M. Scherer, "Schumpeter and the Micro-Foundations of Endogenous Growth," in Horst Hanusch and Andreas Pyka, eds., The Elgar Companion to Neo-Schumpeterian Economics (Edward Elgar: 2007).

$5,000 per project, the payoff distribution described in the previous paragraph would present 24 attractive targets. A precise single-shot "hit" on each of these would yield summed quasi-rents of $346,070, from which R&D costs of 24 x 5,000 = $120,000 must be subtracted to yield a net profit of $226,070.

If under uncertainty the R&D dart throwers are unable to aim with such perfect precision, but instead strew their thrusts randomly over the dart board's payoff space, a quite different outcome ensues. This process was simulated by assuming that each R&D project or dart throw landed at some random coordinate in payoff (dartboard) space. The distribution of possible coordinate "hits" was assumed to be uniform random with replacement. The simulation was performed across an array of eight different sample sizes (i.e., number of dart throws), and for each such sample, the procedure was replicated ten times with a fresh sample of randomly determined coordinate "hit" locations. Multiple hits on a single coordinate location were assumed to add no quasi-rent value, e.g., as if the payoffs from a double hit were shared between two firms marketing a product with identical characteristics. Since all trials were in effect carried out before results were tallied, the approach is analogous to the pure single-stage parallel paths strategy analyzed in earlier sections, with no series strategy component.

With R&D costs of $5,000 per dart throw and assuming 24 throws, as in the perfect-aim case above, the average quasi-rent sum per iteration across ten 24-throw iterations was $137,020. Subtracting from this average payoff total R&D costs of 24 x 5,000 = $120,000, the average net profit per iteration was $17,030. Thus, the innovative efforts yielded on average only slight supra-normal profits -- lower proportionately than those found for most of the parallel and series strategies analyzed earlier. However, as one expects in sampling from highly skew potential payoff distributions, the results of the ten 24-shot iterations varied widely, with net profits after the deduction of R&D costs ranging from $146,750 down to -$73,010.

The strategy pursued by dart-throwing R&D managers depends not only upon the distribution of potential payoffs but on the cost per R&D project (dart throw), which together determine the optimal number of throws. Figure 9 illustrates the dependency of net profits (i.e., summed quasi-rents less total R&D costs) upon the average cost per R&D project. With zero cost per R&D project (the top solid line in Figure 9), one continues throwing darts beyond 100 darts per iteration in the hope of hitting previously untouched payoff coordinates.[21]

---

21 . There can be numerous multiple hits (ex post, duplicative R&D projects). With 100 trials

However, when each dart throw entails cost, the attractiveness of large-numbers attacks eventually diminishes, the more so, the higher the cost per throw. With R&D cost of $1,000 per dart throw, an extension of the experiments recorded in Figure 5 shows, it is profitable to increase the number of throws per iteration to at least 200. A similar extension reveals that with an R&D cost of $2,000 per throw, the net income-maximizing number of throws is on the order of 100. With costs of $3,000 per throw, the profit maximum lies in the range of 24 to 40 throws. (It is impossible to be more precise because, with such a highly skew payoff distribution, considerable variability of outcomes remains even after the experiment is iterated ten times.) With even higher R&D costs per throw, the profit maxima lie in the range of 10 to 24 throws, although again, the intrinsic variability of the results precludes greater precision. Quite generally, and completely consistent with results from the models of optimal parallel paths strategy presented earlier, the greater is the surplus of average payoffs over R&D costs for any given number of trials, the larger is the profit-maximizing number of trials.

When average payoffs are large relative to R&D costs, undertaking numerous trials (dart tosses) may be attractive not only because total payoffs rise by more than R&D costs, but also because proliferation of attacks on skew-distributed payoff targets reduces the relative variability of project outcomes. For the experiments described here, the coefficient of variation (i.e., the ratio of the standard deviation of payoffs to the average payoff) for repeated dart-tossing iterations ranged from 0.79 (for five tosses per iteration) down to 0.20 (for 100 tosses). This hedging benefit is characteristically sought by the managers of high-technology venture fund portfolios. On average, the typical U.S. venture fund invests in roughly 40 individual start-up enterprises. Beyond this, the costs of overseeing and managing the diverse investment targets tend to increase disproportionately, discouraging further portfolio diversification.

6. Conclusion

That uncertainty is an important feature of research and development is a

---

or darts, the number of "missed" payoff cells ranged from 35 to 40, implying at least that number of multiple hits.

For the experiments with smaller samples, there was substantial outcome variability among the experiments. The initial result with samples of 20 was particularly low, kinking the lines in Figure 9 downward. Three additional runs of 10 experiments each were conducted for the n = 20 case, with the average result substituted in the version of Figure 9 presented here.

truism. The uncertainties are of two main types: technological, i.e., whether a particular approach "works," and demand-driven, i.e., how consumers respond to the technical solutions achieved. For both kinds of uncertainties, parallel paths strategies are an important coping approach. They may be adopted by a single firm or government agency seeking to meet a market need with new technology, or by the market, i.e., when numerous firms more or less simultaneously pursue their own approaches to meeting a perceived market need. In either case, the analyses above yield some strong clues as to effective strategies. Most importantly, the higher the value of individual successes for a given quantum of uncertainty and cost per trial, the more parallel paths should be pursued.[22] And the greater the uncertainty for a given solution value -- i.e., the lower the probability of single-trial success or the more skew the distribution of market value outcomes -- the more parallel paths one should optimally pursue. The experiments reported in this paper suggest that pure parallel paths strategies are often not optimal in their own right, especially when expected payoffs contingent upon success are modest. Then some combination of parallel and series strategies is likely to be warranted, especially when researchers can learn from their failures and when some approaches are considered more likely ex ante to succeed than others. The quantitative experiments reported in this paper do not yield specific solutions for individual R&D decision-making situations. However, they point to the kinds of strategy options R&D managers should evaluate as they pursue their important work.

---

22 . When social benefits exceed the private benefits appropriable by innovators, as is commonly the case, a larger number of parallel paths is socially optimal than is profit-maximizing for individual market participants, although competition to be a first mover may drive the two closer. See Scherer, "Pharmaceutical Innovation," supra note 14.

FIGURE 1

The Development Possibility Curve for a Weapon System of a Given Quality Level

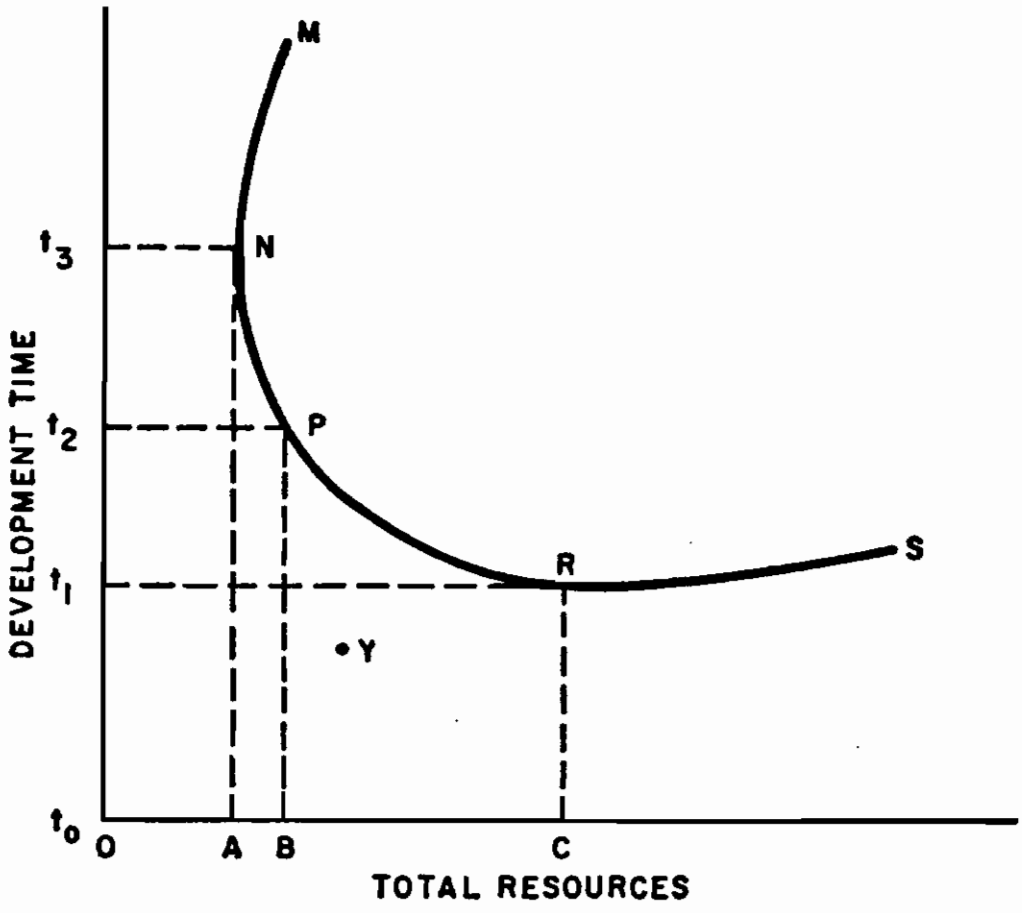

Figure 2

Time-Cost Tradeoffs in Nelson's Original Analysis

(Review of Economics and Statistics, November 1961)

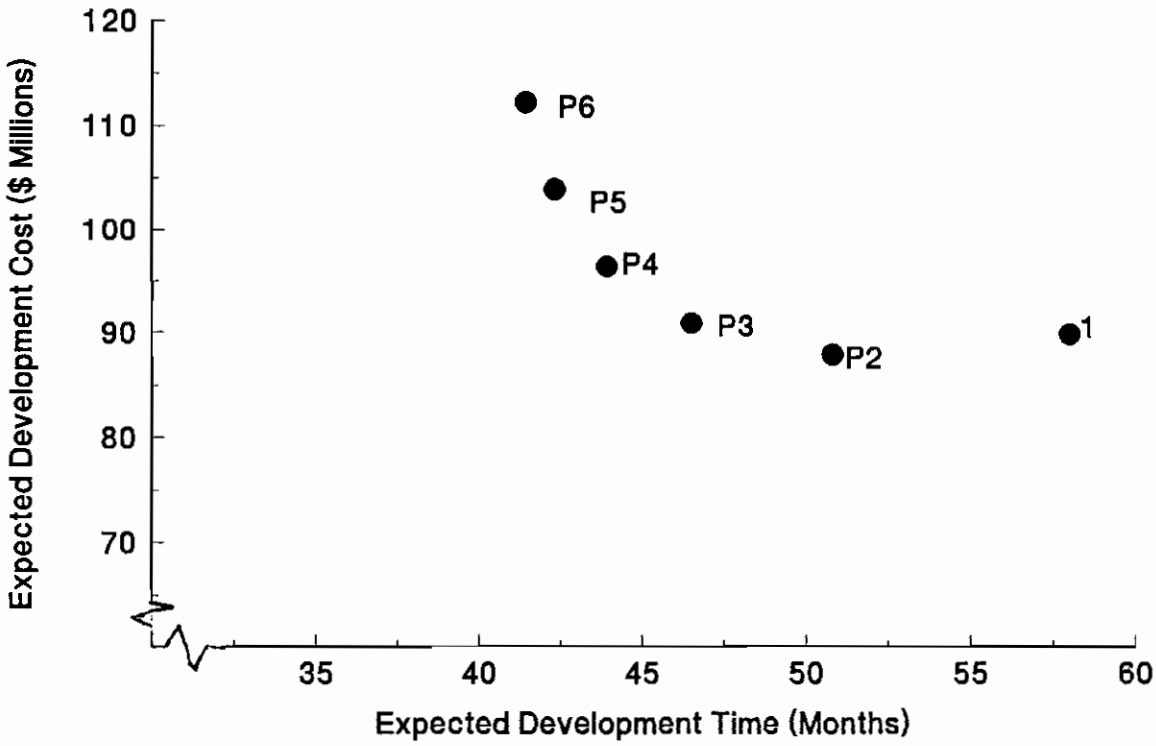

Figure 3

Parallel vs. Series Approaches with Nelson's Assumptions

(Review of Economics and Statistics, November 1961)

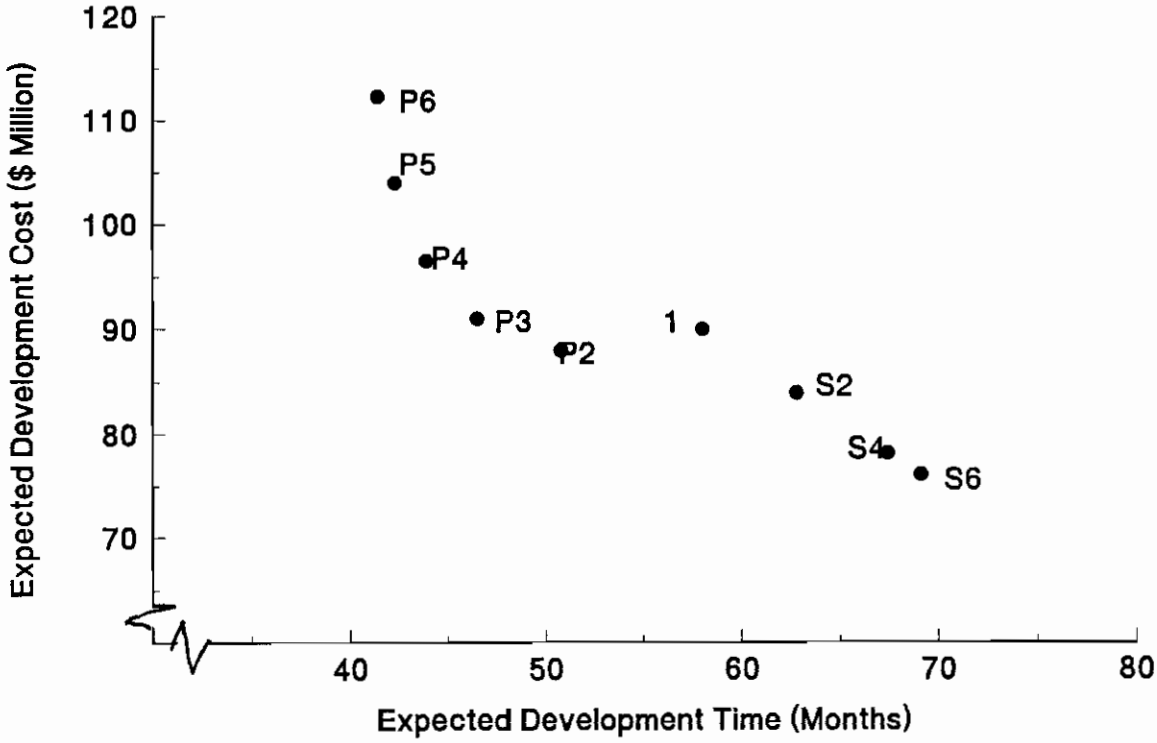

Figure 4
Optimal Choice of Parallel Paths Strategies

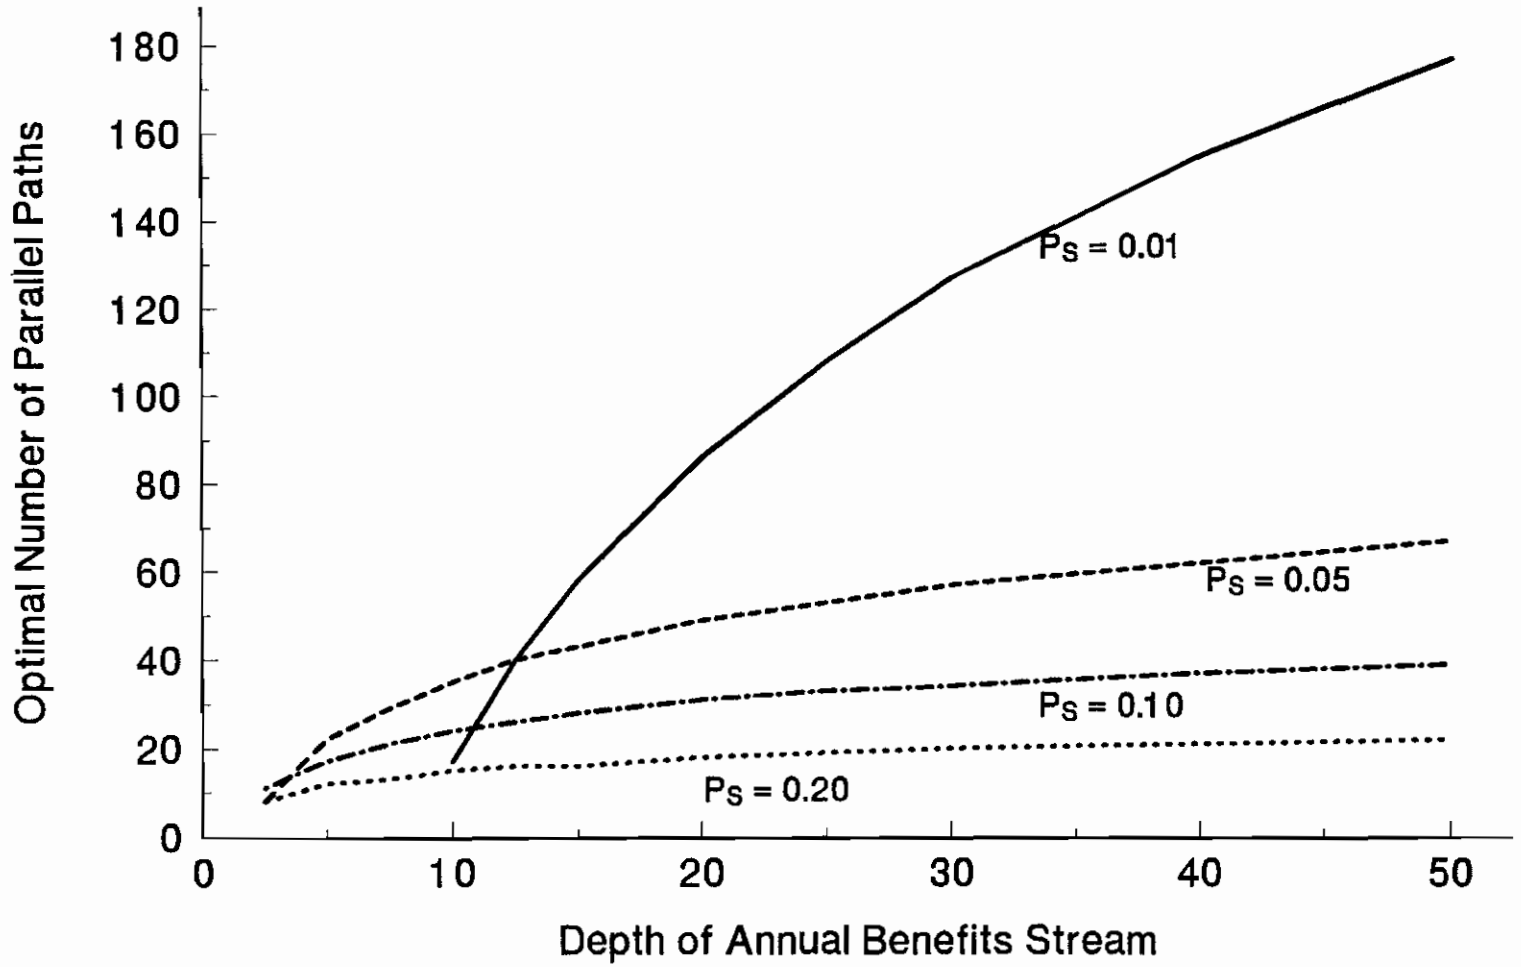

Figure 5

Optimal Trials per Period with Series Strategy

Probability of Single Trial Success = 0.01

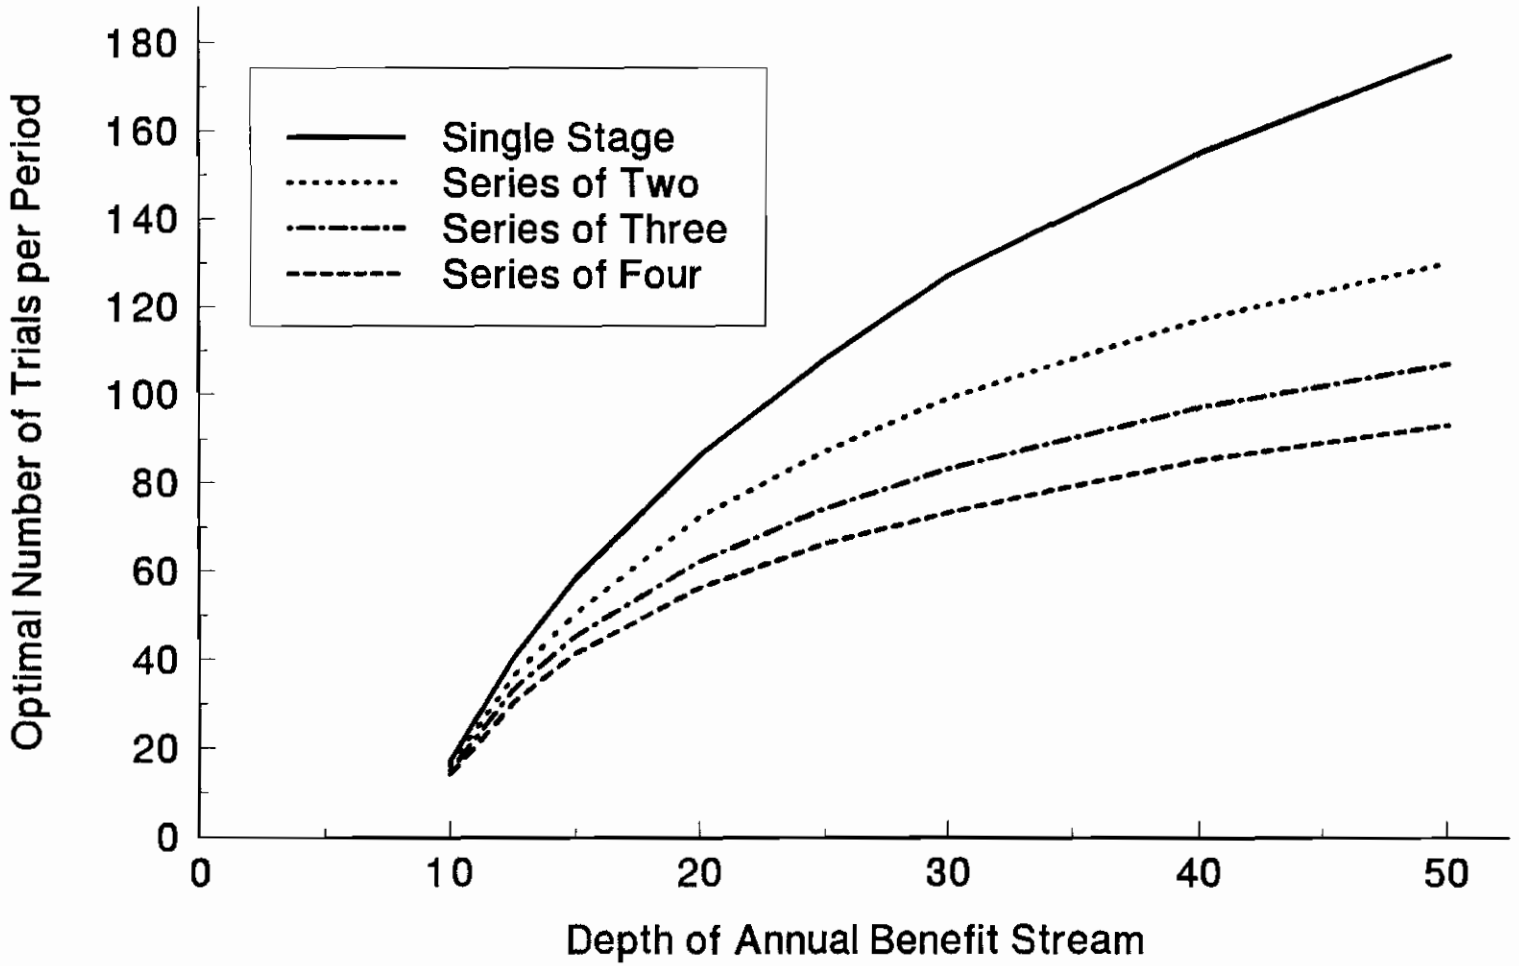

Figure 6

Profits from Optimal Parallel - Series Combinations

Probability of Single-Trial Success = 0.01

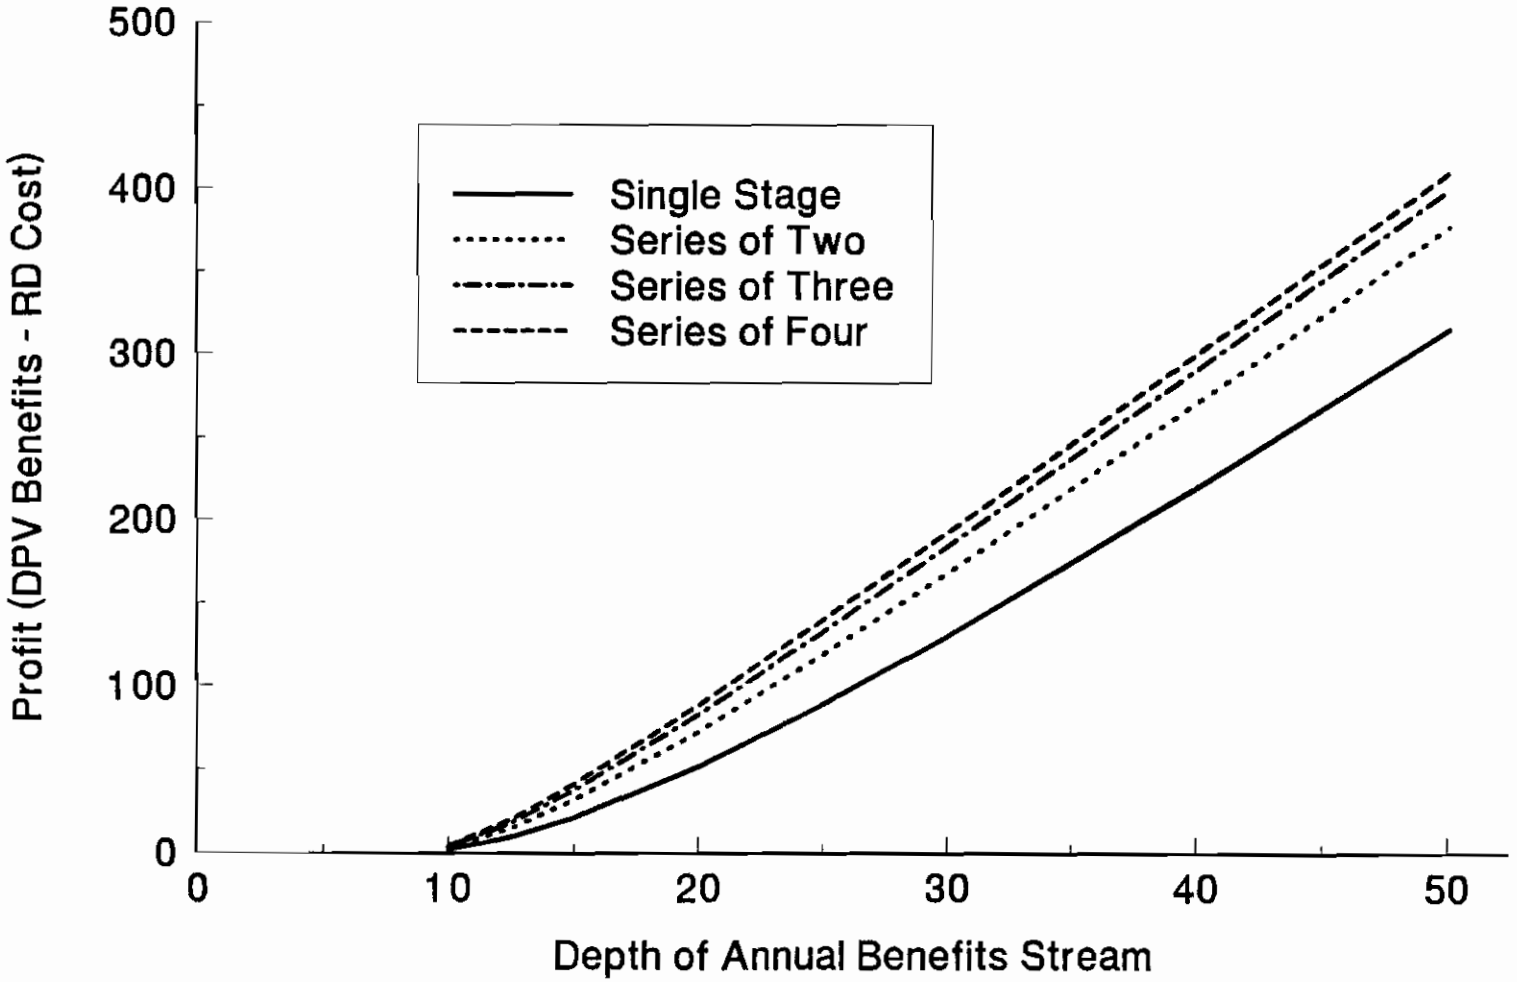

Figure 7

Profits from Using Optimal Two-Stage Number of Trials

Probability of Single-Trial Success = 0.01

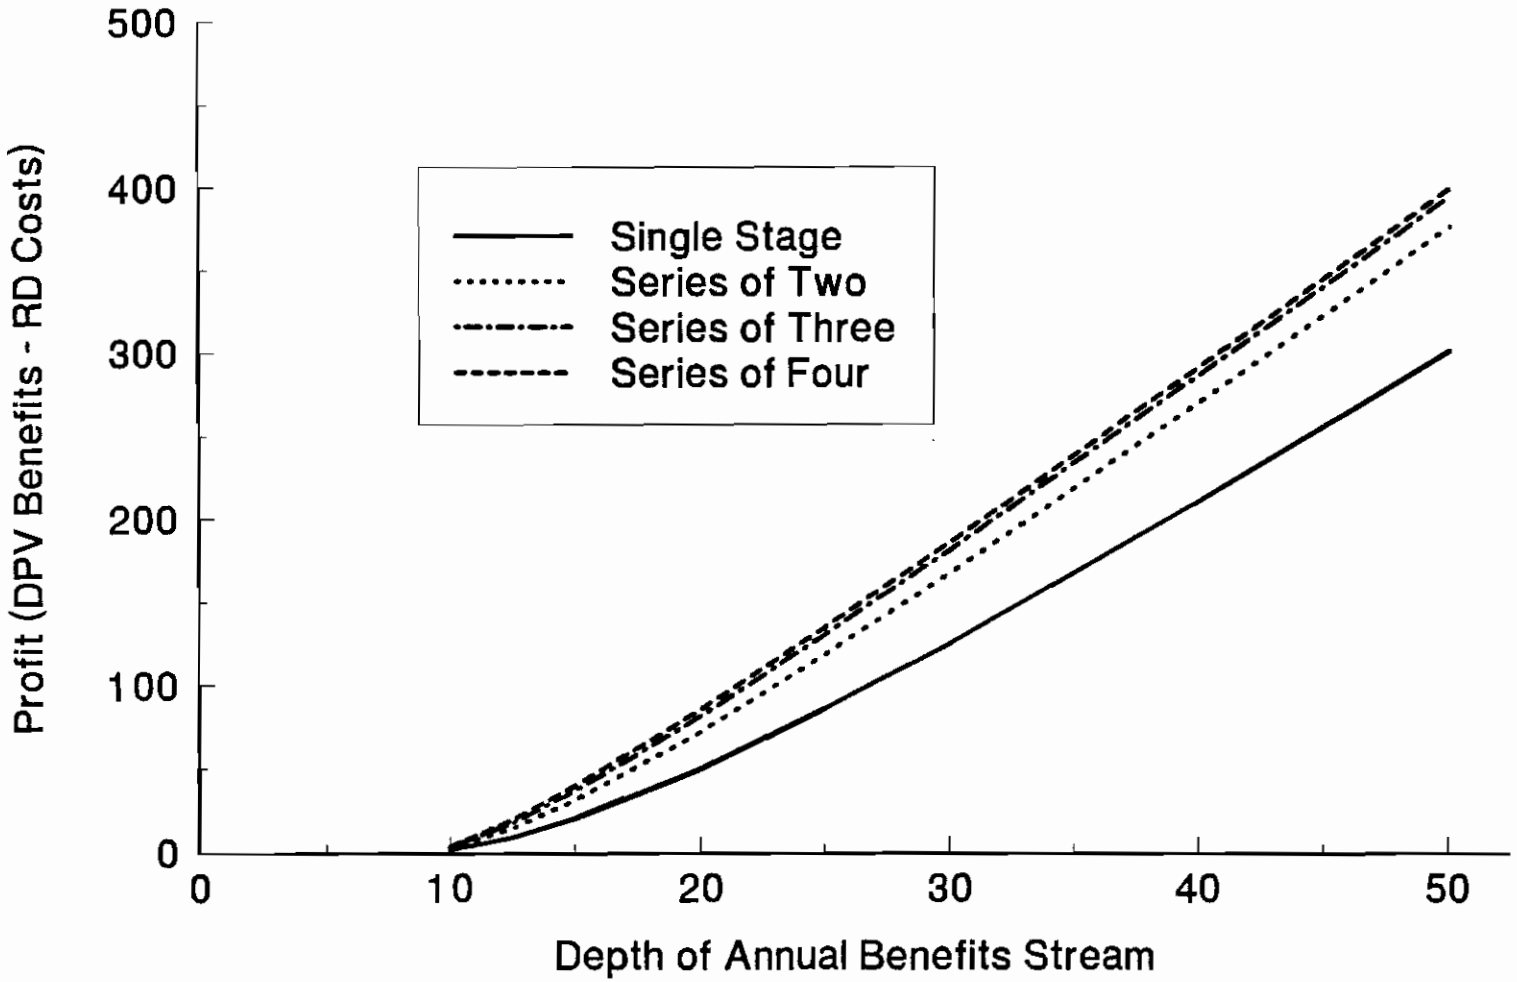

Figure 8

Profits from Single-Year Parallel Paths Strategies

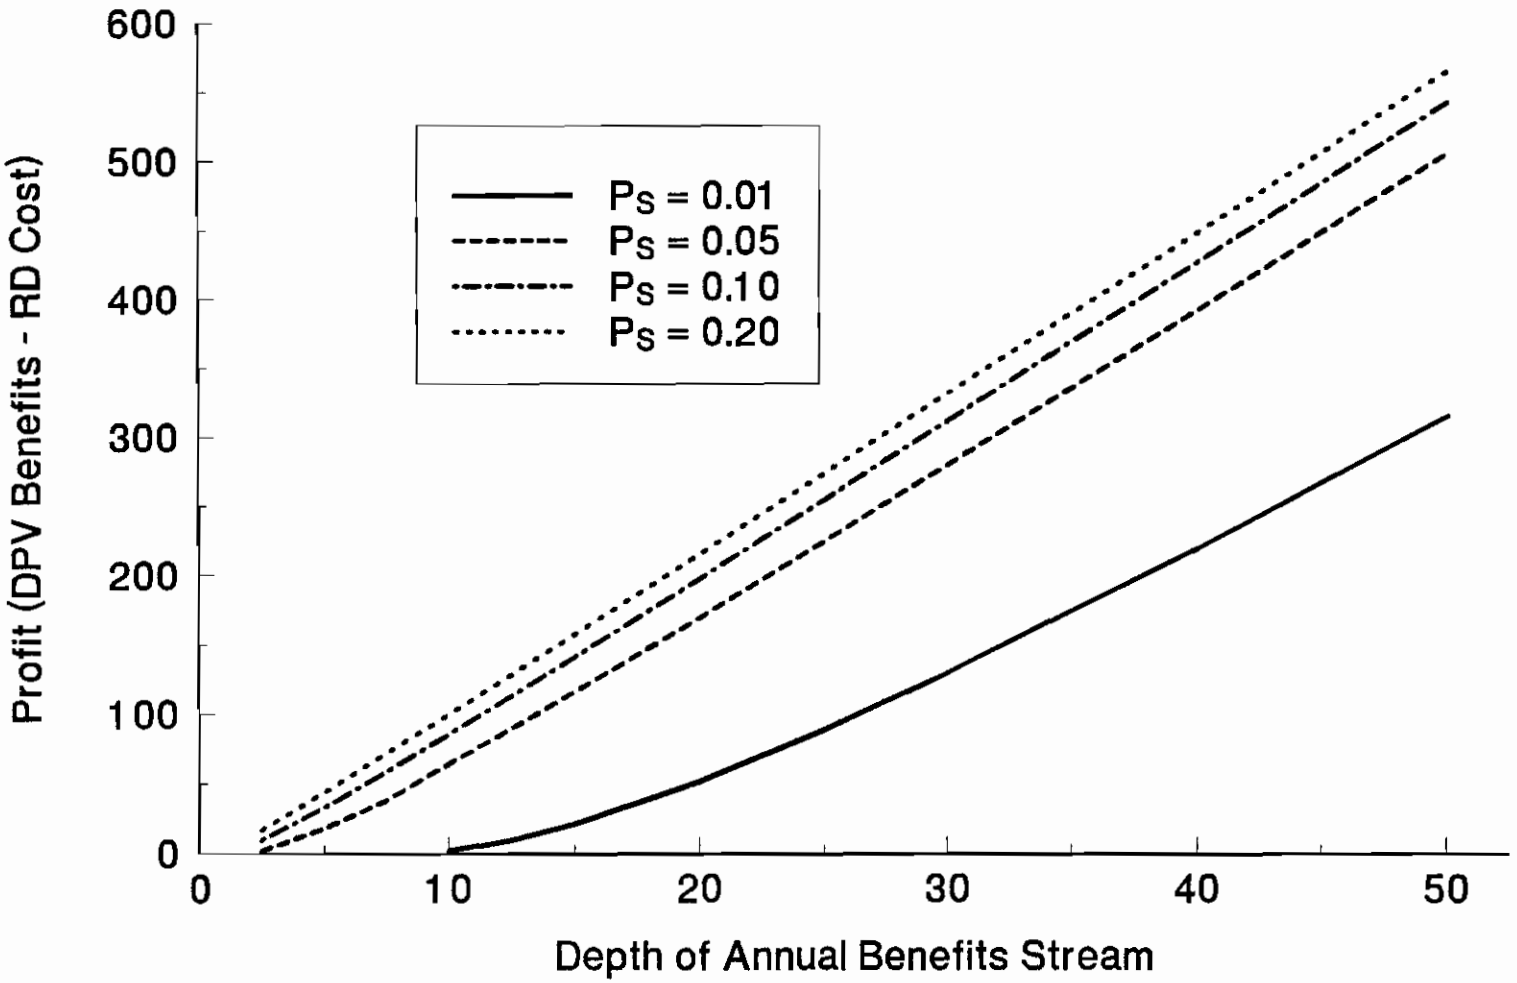

Figure 9

Net Payoffs from Dartboard Experiment